

Usage Analysis of a Large Public Wireless LAN

Timo Ojala, Toni Hakanen, Timo Mäkinen and Veikko Rivinoja

MediaTeam Oulu
University of Oulu, Finland
{firstname.lastname}@ee.oulu.fi

Abstract—This paper presents a quantitative analysis of the usage of a large public wireless local area network which comprises of several components providing outdoor coverage in a city center and indoor coverage at a university campus and at selected public city premises. Analysis is based on correlating two types of data, RADIUS authentication data for determining the true usage of the network, and access point syslog data for localizing the usage. Analysis reveals interesting differences in the daily and weekly usage patterns of different user groups. Analysis of user mobility shows that users are not very mobile, as less than 10% of the sessions may involve spatial movement of at least 50 meters. 60% of the stations using the network appear to have a “home” access point, where the user spends about 80% of the time.

Wireless communication; network monitoring; user mobility

I. INTRODUCTION

Thanks to license-free, economic and mature technology Wireless Local Area Networks (WLANs) are being deployed at an ever increasing speed. Practically all larger academic institutions and corporations have installed a WLAN network on their campuses. Commercial operators have installed tens of thousands of WLAN hotspots at hotels and airports. Many municipal authorities have already deployed or at least have announced plans to deploy city networks [8][14]. Large metropolises are joining the frenzy, for example the City of Philadelphia generated a plenty of headlines in September 2004, by exposing their initiative of investing up to 10 million USD to provide WLAN coverage over the city’s 135 square miles [16].

The wide deployment of WLAN networks is partially motivated by the increasing range of WLAN compatible user devices. Most new laptops contain IEEE 802.11 adapter, but also new WLAN capable devices are entering the consumer market. For example, the new Nokia 9500 Communicator boasts IEEE 802.11b radio in addition to the conventional mobile phone radio.

The increasing availability of both WLAN networks and devices paves way to new types of services such as VoWLAN (Voice over WLAN) or VoIP (Voice over IP) in more traditional terms, videoconferencing and streaming multimedia. While very attractive concepts, the strict QoS (Quality of Service) requirements of real-time interactive multimedia applications may prove difficult to satisfy in a shared-medium WLAN.

Successful deployment of large WLAN networks, together with new multimedia services, calls for understanding the usage of the WLAN networks and the related services.

Consequently, several studies presenting analyses of corporate or campus WLAN networks have been recently published [1][3][4][6][7][13][15]. The most extensive in terms of scope is probably the work of Henderson *et al.* [6], which analyzed 1706 users across 476 WLAN AP’s (access points) for the duration of 12 weeks.

In the following we present a usage analysis of the panOULU public access WLAN network located in the city of Oulu in Finland. The analysis is based on correlating two types of data, RADIUS (Remote Authentication Dial-In User Service [11]) authentication data for determining the true usage of the network, and syslog data generated by the panOULU AP’s for localizing the usage.

There are few important issues, which distinguish our work from previous studies:

- Multifaceted topology of the network: outdoor network in a city center, a campus indoor network and an indoor network in public city premises. This allows us to compare usage patterns at different environments. Previous studies have been restricted to the analysis of either campus [4][6][7][13], corporate [3][15] or conference networks [1].
- Long duration of the monitoring: 14 months. Previous studies have been restricted to 12 weeks [6][15], 11 weeks [4], two months [7], four weeks [3], a week [13] or few days [1].
- Central RADIUS server, which allows accurate session determination. Only one [13] of the previous studies had a similar setup.

II. PANOULU NETWORK

panOULU is a public wireless access network in Oulu in Northern Finland. panOULU network is based on the cooperation agreement which City of Oulu, University of Oulu, Oulu Polytechnic and Oulun Puhelin Plc., the local telco, signed in October 2003. By the agreement these four actors combined their public visitor networks into the panOULU network, which can be accessed by the account holders of the actors and their visitors via e.g. a WLAN connection.

A. Network structure

Figure 1 illustrates the simplified logical structure of the panOULU network. The network comprises of four WLAN networks: city center outdoor network (RotuuariWLAN, named

after the main walking street of the city), indoor network at selected public premises of the City of Oulu (OukaWLAN, named after the city’s Internet domain *ouka.fi*), and the campus networks of the University of Oulu (KampusWLAN) and the Oulu Polytechnic (OuluNET). Two of the networks, KampusWLAN and OuluNET, are connected to the public Internet.

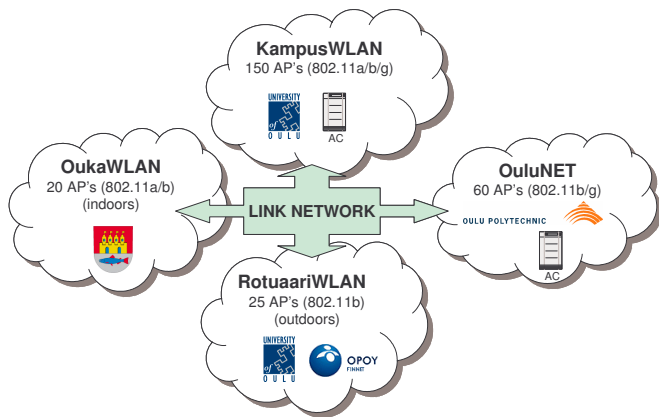


Figure 1. The simplified logical structure of the panOULU network.


RotuaariWLAN contains currently 25 IEEE 802.11b AP’s located as shown in Fig. 2(a). The omni-directional antennas of the AP’s are mounted on building walls at the height of 2-3 stories, mainly at street junctions. The “backhaul” connection of the AP’s is realized with LRE (Long Reach Ethernet) technology. The rather dense placement of the AP’s is partially motivated by the purpose to facilitate positioning of WLAN clients. Given a pre-recorded calibration data, the Ekahau Positioning Engine [5] gives quite reliable location estimates of 5-10 meter accuracy, which can then be utilized to provide location-based services [9][11].

Although not shown on the map of Fig. 2(a), there are at the moment also 20 IEEE 802.11a/b AP’s of the OukaWLAN network at the city center. They provide indoor coverage in selected public premises such as city hall, main library, youth and culture center, and science center.

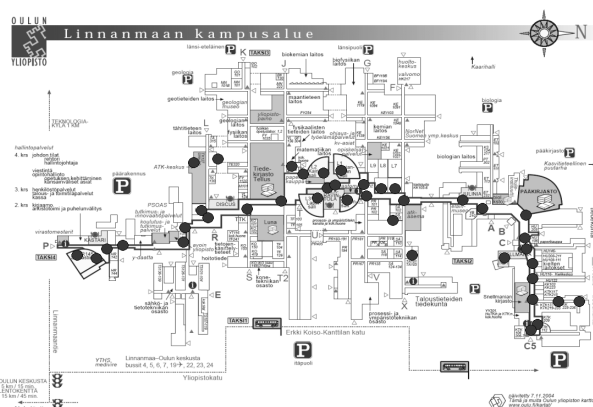
Fig. 2(b) shows the overall layout of the wireless campus network of the University of Oulu, which currently has about 150 IEEE 802.11a/b/g AP’s. Only the locations of selected AP’s providing indoor coverage at key areas such as lecture halls, meeting rooms and cafeterias are shown. The KampusWLAN includes uniform coverage throughout the main passage spanning the main building, which is slightly over a kilometer in length. Again, in addition to providing plain Internet access, the wireless network can be used for realizing location-based services [1].

The fourth component of the panOULU network is the campus network of the Oulu Polytechnic (named OuluNET), which currently comprises of about 60 IEEE 802.11b/g AP’s.

Thus, the panOULU network currently comprises of about 250 AP’s, which make it one of the largest *public and open* WLAN access networks in the world. This work focuses on the usage of KampusWLAN, OukaWLAN and RotuaariWLAN networks, which are currently operated by the University of Oulu. They reside in the same layer 2 network, which facilitates true IP mobility between them. In the network all those IPv4 services can be used, which operate regardless of NAT (Network Address Translation). A panOULU specific SMTP server (*smtp.panoulu.net*) is provided for sending email. IPv6 works without authentication or NAT.



(a)



(b)

Figure 2. Two main components of the panOULU network:
 (a) RotuaariWLAN - outdoor network in the city center;
 (b) KampusWLAN - the campus network of the University of Oulu.

B. Network usage

About 110000 accounts entitle to the free use of the panOULU network: 67000 library card and teacher accounts of the city of Oulu, 18000 University of Oulu staff and student accounts, 9000 Oulu Polytechnic staff and student accounts, and 16000 broadband accounts of Oulun Puhelin Plc. Further, the actors distribute free fixed-term visitor accounts hence panOULU provides open Internet access to everybody.

Authentication is realized with an access controller (AC), which employs RADIUS authentication. To obtain Internet access the user needs to open a web browser and connect to any site. The AC captures the HTTP connection and presents the user with a login window, where the user has to enter his/her username, password and choose the corresponding domain from a pull-down menu. The AC then asks the RADIUS proxy to authenticate the username and the password with the RADIUS server of the designated domain. The AC has a capacity of 1000 concurrent sessions. The lone exception is provided by the “whitelisting” of the URL’s of the websites of the actors providing the panOULU network, i.e. their websites can be browsed without authentication.

Additionally, panOULU is a member of the WLAN roaming pilot of the CSC – Scientific Computing Ltd. (Finnish IT center for science). It basically facilitates accessing the panOULU network with the accounts of other Finnish universities, which are participating in the roaming pilot.

III. USAGE ANALYSIS OF THE PANOULU NETWORK

We have since the inception of the panOULU network collected data for the purpose of analyzing the usage of the panOULU network. Understanding the usage is important for further development of the network and the applications and services offered in the network. We are for example interested in user mobility, as it has great relevance to the network topology.

A. Data

Our study is based on two types of data: RADIUS data collected by the RADIUS proxy, and syslog data generated by the AP’s and stored by a central log server. The data corresponds to RADIUS sessions requested by stations associated to the AP’s in the RotuaariWLAN, the KampusWLAN and the OukaWLAN networks. While these three networks at the moment total about 200 AP’s, the combined size of the network was only half of that at the beginning of the monitoring in January 2004.

The RADIUS log data comprises of three principal types of entries created at the beginning (Acct-Status-Type = Start in RADIUS terminology [8]) and the end (Acct-Status-Type = Stop) of the session, and the keep alive events (Acct-Status-Type = Alive) in between. The Keep-Alive interval of our AC is set to 1200 seconds, i.e. the AC pings each client every 20 minutes. If there is no response, then the AC pings the client 10 times at 5 second intervals. If the client does not respond, then the session is closed. Alternatively, the user can log out explicitly.

Of these three principal types of entries the Stop entries are most useful for our purposes, as they contain complete data of the session, including the MAC address of the WLAN station (Calling-Station-Id), the amounts of input/output data (e.g. Acct-Input-Octets, Acct-Output-Packets), a Unix timestamp denoting the time when the session was closed and the length of the session in seconds (Acct-Session-Time). To remove any spurious entries corresponding to network management, we impose a minimum duration of 10 seconds for the Acct-Session-Time for a Stop entry to be taken into consideration.

We treat the RADIUS Stop entries as the authoritative indication of true network usage. There is a small margin of error there, though, due to the whitelisting of URL’s of the websites of the actors providing the panOULU network.

The APs send timestamped syslog events to a central log server whenever stations associate (ASSOC), disassociate (DISASSOC) or roam (ROAM). While our RADIUS data is complete, our syslog data unfortunately is not for various reasons, including server and AP failures, and syslog UDP packets getting lost in the network.

We have been collecting RADIUS and syslog data since December 2003. In this study we provide statistics for the 14 month long period from January 2004 to February 2005. We cross-correlate the two data sets so that the RADIUS data is used for revealing the true network usage, while the syslog data localizes the usage. In other words, we take into account only those syslog events, which correspond to a RADIUS session, i.e. they occur during the RADIUS session.

It should be emphasized that not all syslog events correspond to actual network usage, but they can for example result from an idle station (dis)associating to nearby AP’s when the station is carried around. In our case about 57% of the syslog events do correspond to an active RADIUS session. This is a very interesting observation, since it implies that using plain syslog data for characterizing network usage can lead to erroneous observations.

Since starting a RADIUS session requires the station first associating with an AP, we also take into account the last ASSOC event before the start of the session, if it occurs during the preceding 24 hours (86400 seconds). Our data shows that some idle stations have remained associated for days or even weeks before a RADIUS session has been started, but for clarity we employ the 24 hour limit.

B. Definitions

To formalize our presentation we adopt the following definitions:

Account – Refers to a unique account of form *username@domain*. A single person may have multiple accounts.

Station - Refers to the MAC address of a wireless NIC (network interface card) of a device, e.g. PC, PDA or a mobile phone. While MAC addresses are unique, several users may share a device, or a person may use multiple devices.

Session – Refers to a session in the RADIUS data. A session is initiated by a station.

Mobile session – A session is deemed to be mobile if the station initiating it is associated with multiple AP’s during the session and the spatial distance between these AP’s exceeds a threshold. This is illustrated in Fig. 3, where station X is associated with three AP’s, A, B and C, during a session. If $\max(AB, BC, CA)$ exceeds a threshold, then this sessions is regarded mobile. Following Henderson *et al.* [6], we also use 50 meters as the spatial threshold in determining mobile sessions.

WLAN technology is known to suffer from the “ping pong” effect, where a station roams back and forth between two adjacent AP’s without any real station mobility taking place. This kind of artificial mobility was excluded by ignoring sessions, where a station roamed between two particular AP’s more than twice.




Figure 3. The session involving station X is deemed to be mobile if $\max(AB, BC, CA)$ exceeds a threshold (50 m in this study).

Home AP - If at least 50% of a station’s associations are to a single AP, then this AP is deemed as the home AP of the station.

C. General usage characteristics


According to the RADIUS data 5171 unique accounts logged into the network 75539 times using 4115 stations during the 14 month period. The larger number of accounts is explained by the fact that certain stations such as those loaned to patrons in the city library are used by multiple persons (accounts).

The 75539 sessions are distributed over different user groups as follows: staff and students of the University of Oulu 73.1%, accounts of the city of Oulu 15.1%, staff and students of Oulu Polytechnic 6.6%, and visitor accounts 2.8%.

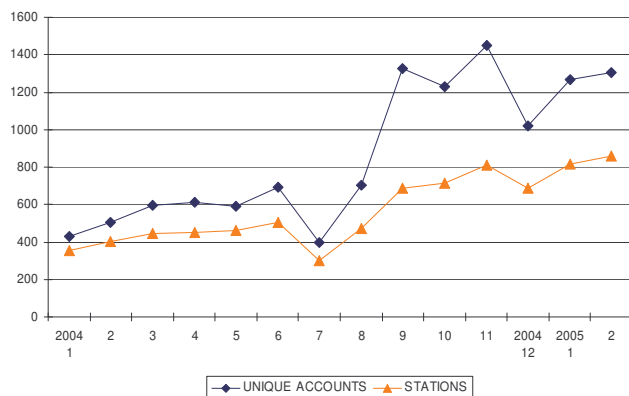
The charts in Fig. 4 illustrate some general usage characteristics over the 14 month period on a monthly interval. The chart in Fig. 4(a) shows that the number of active AP’s involved in RADIUS sessions per month increased from 106 to 170, as new AP’s were deployed. The number of concurrent users ranged between 32 and 115.

The chart in Fig 4(b) shows that the number of unique accounts per month ranged from 431 to 1449, whereas the number of stations ranged from 352 to 857. The dip in the network usage in summer 2004 is due to the summer holidays of university staff and students.

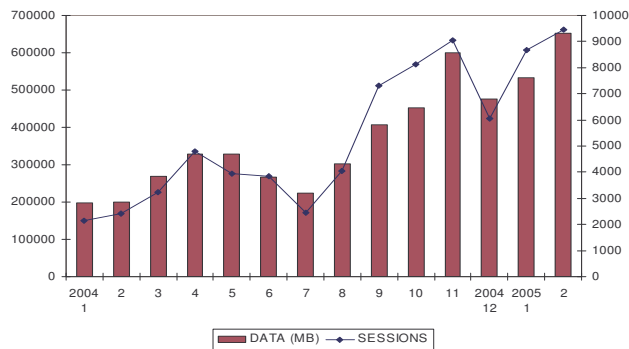
The chart in Fig. 4(c) shows that the number of RADIUS sessions per month has increased from about 2150 in January 2004 to about 9450 in February 2005. At the same time the amount of transferred data has increased from 200 GB in January 2004 to 650 GB in February 2005.



(a)



(b)




(c)


Figure 4. General usage characteristics.

The chart in Fig. 5(a) shows the daily distribution of the starting times of RADIUS sessions of selected user accounts. As expected, most of the sessions are initiated during the office hours with peak activity between 11am and 3pm, whereas the smallest activity takes place during the early morning hours. We can observe interesting differences between different user groups. The use of university and visitor accounts concentrates more clearly on the office hours, while the city accounts see the highest relative activity during the evening hours.

The chart in Fig. 5(b) shows the weekly distribution of the starting times of all RADIUS sessions of selected user accounts. Again, the network usage by university and visitor accounts tends towards office days, whereas city accounts exhibit clearly higher relative activity during weekends, which can be attributed to the home use at the city center.



(a)




(b)

Figure 5. Daily (a) and weekly (b) distribution of the starting times of the RADIUS sessions of selected user accounts.


The chart in Fig. 6(a) shows the median amount of input/output data per session in megabytes in each of the 14 months. We are using median to characterize the typical amount of input/output data per session, because it is a more robust statistic in the presence of few sessions with exceptionally large amounts of traffic which would skew mean, for example. As expected, the inbound traffic is much higher, on average about 6.6 times the outbound traffic.

The charts in Fig. 6(b-c) show that the RADIUS sessions of the city accounts have clearly larger amounts of inbound traffic than the university accounts, which become clearly more active towards weekends, as well. Fig. 6(b) shows an outstanding peak of activity around 2-4am by city accounts. While the overall ratio of inbound and outbound traffic is 6.6, during the

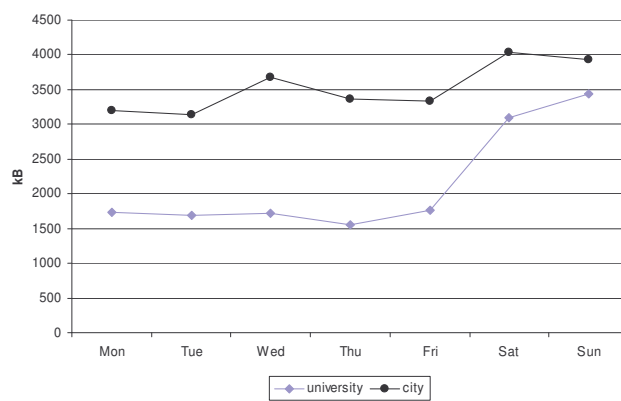
early hours it is over 10 for both user groups. Both these observations, together with the occurrence of sessions with exceptionally large amounts of inbound/outbound traffic (up to 4.0 GB per session), speak for P2P usage.



(a)



(b)

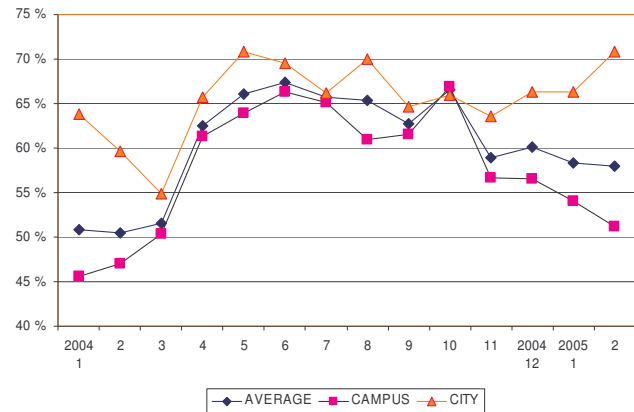


(c)


Figure 6. (a) Median amount of input/output data per session in each month; (b) Median amounts of input data for RADIUS sessions started at the given hour by university and city accounts; (c) Median amounts of input data for RADIUS sessions started on the given day by university and city accounts.

D. User mobility

The charts in Fig. 7 illustrate statistics of stations having a home AP. Fig. 7(a) shows that in the city on average 66% of the stations have homes, whereas the proportion of home stations is clearly smaller on the campus, about 58%. Fig. 7(b) shows the proportion of associations to the home AP of all associations. On average 60% of stations have homes where they spend 81% of time. This is fairly close to the result of Henderson *et al.* [6], who found that 50% of stations had a home where they spent 74% of their time.



(a)



(b)

Figure 7. Statistics on stations having a home AP: (a) the percentage of stations having a home AP of all stations; (b) the percentage of associations to the home AP of all associations.

The chart in Fig. 8 illustrates the proportion of mobile sessions of all sessions. On average 9% of the sessions in the city are mobile, while the corresponding figure in the campus is 5%. It would be daring to draw any conclusions on different mobility patterns, since our crude measure against the “ping pong” effect is not guaranteed to remove all superficial mobility caused by arbitrary roaming between AP’s, which is

much more likely to occur in the open outdoor space of city center than indoors on the university campus.

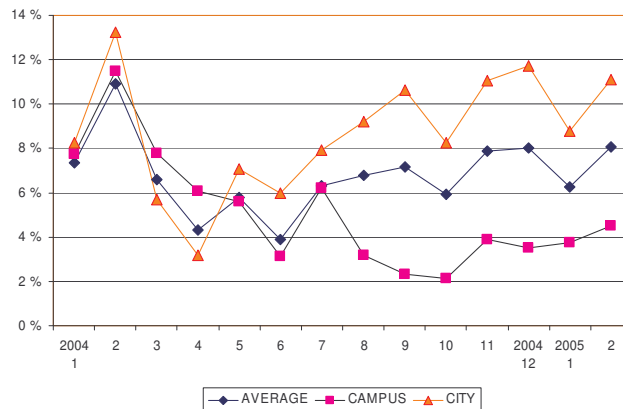


Figure 8. Percentages of mobile sessions of all sessions.

IV. CONCLUSIONS

We presented the usage analysis of a large versatile public WLAN access network over a 14 month period. The analysis was based on correlating two types of data, RADIUS authentication data for determining the true usage of the network, and access point syslog data for localizing the usage. In our case 43% of the syslog events were found not to correspond to any active RADIUS session. This implies that using plain syslog data for characterizing network usage can lead to erroneous observations.

Our analysis revealed interesting differences in the daily and weekly usage patterns of different user groups. While university accounts tend towards office hours and weekdays corresponding to work usage, the city accounts exhibit clearly higher relative activity during the evening hours and weekends corresponding to home usage. In terms of the volume of inbound traffic city accounts are more active than university accounts, as well. Within the daily cycle both user groups have exceptionally high ratio of inbound/outbound traffic during the early morning hours.

A more detailed analysis of user mobility showed that users are not very mobile, as less than 10% of the sessions may involve spatial movement of at least 50 meters. A majority of the stations using the network appear to have a home AP, where the user spends about 80% of the time. The user mobility can be expected to increase once devices such as VoWLAN/VoIP phones become more prominent, as talking to a phone is much more natural than typing the keyboard of a laptop while moving.

We are currently conducting a deeper analysis of the massive data, to be able to provide a more detailed insight to the network usage, including profiling the different types of traffic in the network.

ACKNOWLEDGMENTS

The authors wish to thank the panOULU consortium, the National Technology Agency of Finland, Cisco Systems Finland Ltd., Nokia Plc. and Oulun Liikekeskus ry. for their invaluable support.

REFERENCES

- [1] M. Aittola, T. Ryhänen and T. Ojala, "SmartLibrary - Location-aware mobile library service," Proc. Mobile HCI 03, Udine, Italy, pp. 411-416.
- [2] A. Balachandran, G. M. Voelker, P. Bahl and P. V. Rangan, "Characterizing user behavior and network performance in a public wireless LAN," Proc. ACM SIGMETRICS 2002, Marina Del Rey, CA, 2002, pp. 195-205.
- [3] M. Balazinska and P. Castro, "Characterizing mobility and network usage in a corporate wireless local-area network," Proc. ACM MobiSys 2003, San Francisco, CA, 2003, pp. 303-316.
- [4] F. Chinchilla, M. Lindsey and M. Papadouli, "Analysis of wireless information locality and association patterns in a campus," Proc. IEEE Infocom 2004, Hong Kong, China, 2004, vol. 2, pp. 906-917.
- [5] Ekahau Positioning Engine, Ekahau Inc., <http://www.ekahau.com>, 2005.
- [6] T. Henderson, B. Kotz and I. Abyzov, "The changing usage of a mature campus-wide wireless network," Proc. ACM MobiCom 2004, Philadelphia, PA, 2004, pp. 187-201.
- [7] R. Hutchins and E. W. Zegura, "Measurements from a campus wireless network," Proc. IEEE ICC 2002, New York, NY, 2002, vol. 5, pp. 3161-3167.
- [8] MuniWireless.com, Reports on municipal wireless and broadband projects, <http://www.muniwireless.com>, 2005.
- [9] T. Ojala T, J. Korhonen J, M. Aittola, M. Ollila, T. Koivumäki, J. Tähtinen and H. Karjaluo, "SmartRotuaari - Context-aware mobile multimedia services," Proc. MUM 2003, Norrköping, Sweden, 2003, pp. 9-18.
- [10] panOULU network, <http://www.panoulu.net>, 2005.
- [11] J. Peltonen, M. Ollila and T. Ojala, "TimeMachine Oulu - Dynamic creation of cultural-spatio-temporal models as a mobile service," Proc. Mobile HCI 03, Udine, Italy, 2003, pp. 342 - 346.
- [12] Remote Authentication Dial In User Service (RADIUS), RFC 2865, 2005, <http://www.freeradius.org/rfc/rfc2865.html>.
- [13] D. Schwab and R. Bunt, "Characterizing the use of a campus wireless network," Proc. IEEE Infocom 2004, Hong Kong, China, 2004, vol. 2, pp. 862-870.
- [14] S. A. Shamp, "WiFi clouds and zones: A survey of municipal wireless initiatives," University of Georgia, 2004, <http://www.muniwireless.com/reports/docs/Wificloudszones.pdf>.
- [15] D. Tang and M. Baker, "Analysis of a local-area wireless network," Proc. ACM MobiCom 2000, Boston, MA, pp. 1-10.
- [16] Wireless Philadelphia, <http://www.phila.gov/wireless/>, 2005.